arXiv:1610.03714v2 [quant-ph] 3 Nov 2016

Quantum state estimation when qubits are lost: A no-data-left-behind approach

Brian P. Williams and Pavel Lougovski

Quantum Information Science Group, Oak Ridge National Laboratory, Oak Ridge, Tennessee USA 37831

E-mail: williamsbp@ornl.gov; lougovskip@ornl.gov

Abstract. We present an approach to Bayesian mean estimation of quantum states using hyperspherical parametrization and an experiment-specific likelihood which allows utilization of all available data, even when qubits are lost. With this method, we report the first closed-form Bayesian mean estimate for the ideal single qubit. Due to computational constraints, we utilize numerical sampling to determine the Bayesian mean estimate for a photonic two-qubit experiment in which our novel analysis reduces burdens associated with experimental asymmetries and inefficiencies. This method can be applied to quantum states of any dimension and experimental complexity.‡

Keywords:quantum state estimation, qubit, Bayesian mean estimation

Contents

1 Introduction **2**

2 Approach Outline **4**

3 An example: Bayesian mean estimation of an ideal single qubit **6**

4 Bayesian mean estimation for multi-qubit experiments **11**
- 4.1 Estimating parameters in a single-basis two-photon experiment 11
- 4.2 Single-basis simulation 15
- 4.3 Parametrizing the n-dimensional density matrix 16
- 4.4 Estimating parameters in a multi-basis two-photon experiment 17
- 4.5 State certainty 18

5 BME performance with numerical sampling **19**

‡ This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan(http://energy.gov/downloads/doe-public-access-plan).

6 Experimental tomography **21**

7 Conclusions **22**

8 Acknowledgement **23**

9 References **23**

Appendices **24**

A Maximum likelihood estimation with a traditional likelihood **24**

B Numerically sampling density matrices **24**
B.1 Slice sampling 25
B.2 Density matrix sampling 26

C Maximization using gradient ascent **27**

1. Introduction

The problem of estimating the state of a quantum system from observed measurement outcomes has been around since the dawn of quantum mechanics. In recent years, driven by technological advances in building and controlling quantum systems, this question has received a renewed interest especially in the context of quantum information processing (QIP). Since QIP relies on one's ability to prepare arbitrary multi-qubit quantum states, verifying experimentally that a quantum state has been indeed prepared acceptably close (by some metric) to a target is of paramount importance. As a result, methods built on classical statistical parameter estimation procedures have been adopted in order to satisfy a demand for quantum state characterization tools.

Arguably the most widespread quantum state estimation approach relies on classical maximum likelihood estimation (MLE) technique. Pioneered in [1], it offers a great simplicity in numerical implementation but suffers from several dangerous flaws. Most prominently, MLE is prone to rank-deficient estimates of a density matrix from a *finite* number of measurement samples [2] which in turn implies that probability of observing certain states is zero – a statement that is only valid in the limit of *infinite* number of observations. Also, MLE does not provide a straightforward way to place error bars on an estimated quantum state.

Fortunately, there is an alternative parameter estimation technique called the Bayesian mean estimate (BME) that can be applied to quantum state characterization [2, 3, 4] and is free of the the aforementioned shortcomings. In addition, BME minimizes the mean square error(MSE) [2, 5] i.e. the average square difference between the parameter and its estimate. Thus, BME offers a more accurate estimate of a quantum state. But, BME poses an implementation challenge. It computes a posterior distribution over quantum states for given measurement data by using Bayes'

rule which in turn requires one to calculate the probability of the data by integrating over the manifold of all physical quantum states. While it may be possible to carry out the multi-dimensional integration analytically, ultimate evaluation still may be computationally prohibitive. Thus, numerical routines that use Monte Carlo (MC) methods are applied in order to sample from the posterior distribution. There is a trade-off between the speed and accuracy of BME depending on which MC algorithm is used. To our knowledge two types of MC algorithms were proposed for quantum state BME so far. The first one is the Metropolis-Hastings [6] (MH) algorithm–an example of Markov Chain MC (MCMC)–was adopted in [2]. The second one is sequential MC (SMC) [7]–an importance sampling based algorithm–recently used for adaptive quantum state tomography [8]. The MH algorithm is known for its ability to reproduce probability distributions very accurately at the expense of slow convergence. On the other hand, the SMC algorithm is fast but may converge to a sample that does not faithfully represent the distribution of interest.

When to apply or the BME depends on many factors. For instance, for a small measurement data set and a large number of unknown parameters–a typical situation for multi-qubit systems–the BME is superior to as we demonstrate in Section 5 of this paper. But perhaps even more crucially, the applicability of the BME approach depends on the choice of the form and parametrization of the likelihood function. The experimental likelihood most used in application is a simple multinomial that connects the observed data set directly with the quantum probabilities using Born's rule [1]. This approach assumes the data set, observed measurement outcomes, result directly and only from the quantum state and unitary operations. This is not always the case, since the measurement apparatus often introduces operation bias (not always unitary) and inefficiencies that modify the probability of an experimental observation. Previously, James et al. [9] accounted for bias in qubit operations in their two-photon tomography method using MLE. More recent MLE works such as Gate Set Tomography [10, 11] assume nothing about the qubit state or operations, gates, other than their dimension. However, previous methods stop short of full accounting for non-unitary operations such as qubit loss. Thus, experimentalists may find themselves applying normalizing constants to account for deficiencies in the defined likelihood. These normalizing constants require preliminary experiments to obtain, and the method for obtaining these constants is often not well defined nor reported.

In this paper we develop a BME-based quantum state reconstruction method that utilizes the slice sampling [12] (SS) algorithm which has the accuracy of the MH algorithm but demonstrates faster convergence [13] and is more resilient for a numerical implementation. We show that by using the hyperspherical parameterization of the manifold of density matrices the BME of a state of a single qubit can be computed analytically by using a uniform prior. For a two-qubit system, in a situation when individual qubits may be lost during the measurement process, we apply SS algorithm to the same parameterization and an experiment-specific likelihood demonstrating a computationally stable and efficient way of sampling from the posterior distribution over

the density matrices. We compare the resulting BME estimates to the corresponding MLE estimates as a function of the number of measurements and observe the superiority of the BME method, especially in the limit of small sample sizes.

We begin this paper with a quick outline of our method in Section 2. We derive a closed BME for the ideal single qubit experiment in Section 3. This approachable example illustrates our method and contrasts it with traditional MLE methods. It may also inspire further research into closed-form BME solutions of higher dimensional quantum systems. Next, in Section 4, we derive a likelihood for a finite data two-photon experiment where detector inefficiencies and experimental asymmetries are taken into account. Utilization of this approach results in the real world benefit of eliminating the need to perform preliminary experiments to determine normalization constants. Subsequently, in Section 5 we simulate a multitude of two-qubit photon experiments generating data sets from which we compare the performance of various MLE and BME approaches. Lastly, we apply our estimation to a real world two-photon experiment in Section 6. In the Appendices, we describe a common MLE approach using a traditional likelihood, we detail numerical procedures for sampling density matrices from the true state distribution using slice sampling, and describe the optimization method used in likelihood maximization.

2. Approach Outline

The components of our quantum state estimation pipeline are outlined in Fig. 1. First, we define a model of our experiment by enumerating all the possible outcomes. This enumeration allows us to specify an experiment-specific likelihood $P(\mathcal{D}|\alpha)$, the probability of observing a specific data set $\mathcal{D}$ given the experimental parameters $\alpha = \{\alpha_1, \cdots, \alpha_N\}$. In our case parameters α are elements of a density matrix ρ representing the quantum state to be estimated. Bayes' rule,

$$P(\alpha|\mathcal{D}) = \frac{P(\mathcal{D}|\alpha)P(\alpha)}{P(\mathcal{D})} \tag{1}$$

then allows us to express $P(\alpha|\mathcal{D})$, a posterior distribution (PD) for the variables α, given an observed data set $\mathcal{D}$ and a prior probability distribution $P(\alpha)$.

Next, the BME for a specific parameter α_i given a data set $\mathcal{D}$ is

$$\overline{\alpha_i} = \int d\alpha P(\alpha|\mathcal{D}) \times \alpha_i. \tag{2}$$

We expand our analysis to quantum systems by assuming that parameters α are entries of a density matrix ρ ($\rho_{ij} = \alpha_k$) describing a valid quantum state (i.e. $\rho \geq 0$, $\rho = \rho^\dagger$, $\mathrm{Tr}(\rho) = 1$). Therefore, α's are not independent as we must enforce quantum constraints. To achieve this in a computationally tractable fashion, instead of using the Cartesian parameterization given by α's we parametrize a density matrix ρ utilizing a Cholesky decomposition (see panel **c** in Fig.(1)) and hyperspherical parameters as

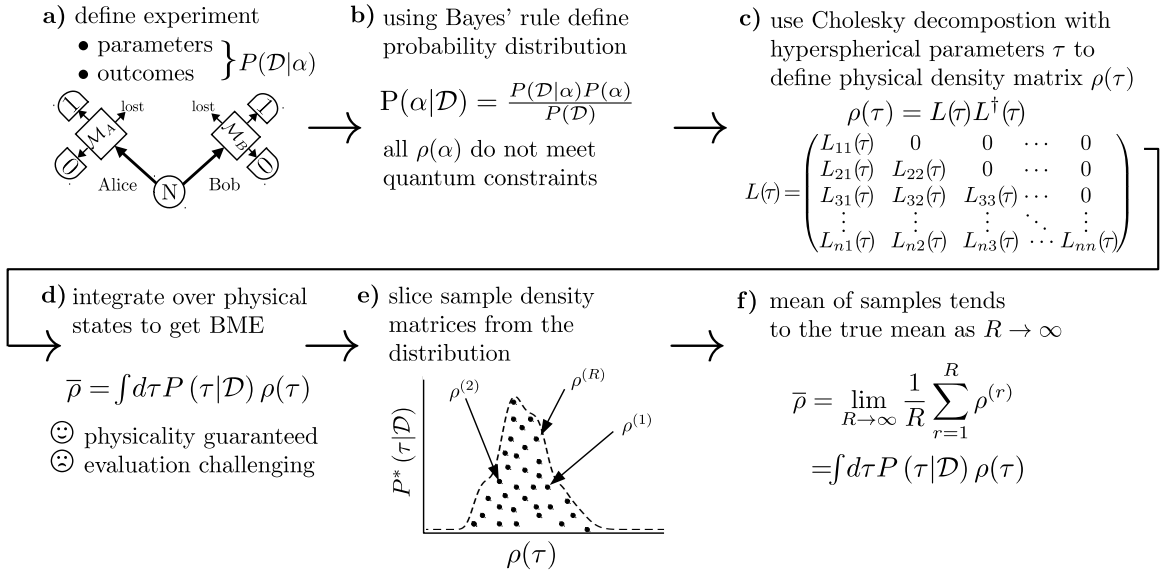
a) define experiment
• parameters
• outcomes
$P(\mathcal{D}|\alpha)$
lost
lost
M_A
M_B
Alice
N
Bob
b) using Bayes' rule define probability distribution
$P(\alpha|\mathcal{D}) = \frac{P(\mathcal{D}|\alpha)P(\alpha)}{P(\mathcal{D})}$
all $\rho(\alpha)$ do not meet quantum constraints
c) use Cholesky decompostion with hyperspherical parameters τ to define physical density matrix $\rho(\tau)$
$\rho(\tau) = L(\tau)L^\dagger(\tau)$
$L(\tau) = \begin{pmatrix} L_{11}(\tau) & 0 & 0 & \cdots & 0 \\ L_{21}(\tau) & L_{22}(\tau) & 0 & \cdots & 0 \\ L_{31}(\tau) & L_{32}(\tau) & L_{33}(\tau) & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ L_{n1}(\tau) & L_{n2}(\tau) & L_{n3}(\tau) & \cdots & L_{nn}(\tau) \end{pmatrix}$
d) integrate over physical states to get BME
$\overline{\rho} = \int d\tau P(\tau|\mathcal{D})\, \rho(\tau)$
physicality guaranteed
evaluation challenging
e) slice sample density matrices from the distribution
$\rho^{(2)}$
$\rho^{(R)}$
$\rho^{(1)}$
$P^*(\tau|\mathcal{D})$
$\rho(\tau)$
f) mean of samples tends to the true mean as $R \to \infty$
$\overline{\rho} = \lim_{R\to\infty} \frac{1}{R} \sum_{r=1}^{R} \rho^{(r)}$
$= \int d\tau P(\tau|\mathcal{D})\, \rho(\tau)$

Figure 1. Our Bayesian mean estimation of density matrices is outlined above. a) We define our experiment by specifying a likelihood function $P(\mathcal{D}|\alpha)$ of data $\mathcal{D}$ given parameters α and b) express a corresponding posterior distribution of the parameters α which define the density matrix given data using Bayes' rule. c) We parametrize the density matrix such that any choice of parameters in a specified range leads to a valid physical state. d) We represent our posterior distribution using these new parameters; now the BME of the density matrix can be formally written down as an integral using a Haar-invariant measure. Unfortunately, the integral's analytical solution is typically computationally intensive to evaluate. e) Thus, we use computationally efficient numerical slice sampling to make samples of ρ from the postirior distribution $P(\tau|\mathcal{D})$. f) As $R\to\infty$ we tend to the true mean $\overline{\rho}$.

suggested by Daboul [14]. We abbreviate this parametrization with τ to distinguish it from the Cartesian parametrization α.

Next, in order to compute the BME estimator, we need to select a prior probability distribution over density matrices $P(\tau) \equiv P(\rho(\tau))$ and an integration measure over the set of all physical quantum states $d\tau$ such that $d\mu(\rho(\tau)) = P(\tau)d\tau$ is a valid probability measure i.e. $\int d\mu(\rho(\tau)) = 1$. We use a non-informative prior $P(\tau) = \text{const}$ and derive the integration measure $d\tau$ induced by the Riemanian metric g_{ij} computed from the Euclidian length element between density matrices $(ds)^2 = \text{Tr}\left(d\rho(\tau) \cdot d\rho^\dagger(\tau)\right)$ [15]. This choice of the integration measure guarantees Haar invariance of the probability measure over the set of density matrices. Thus, the probability of a state $\rho(\tau)$ is invariant under an arbitrary unitary rotation U i.e. $P(\rho(\tau)) = P(U\rho(\tau)U^\dagger)$. Then the BME of an unknown quantum state reads,

$$\overline{\rho} = \int d\tau P(\tau|\mathcal{D}) \times \rho(\tau). \tag{3}$$

The latter expression for the BME can, in principle, be evaluated analytically. However, in practice it almost surely requires computational resources and (or) time

constraints that prohibit analytical evaluation. In this case, an estimate can be obtained using numerical sampling from the posterior distribution $P(\tau|\mathcal{D})$. For example, in later sections we utilize numerical slice sampling [12] to arrive at approximate estimates for a two-photon experiment.

3. An example: Bayesian mean estimation of an ideal single qubit

Consider an ideal single-qubit experiment. In this experiment we can reliably and repetitively prepare a qubit in an unknown state ρ and measure the value of any desired observable M without err or qubit loss. If M is a two outcome POVM defined by operators M_i with $i \in \{0,1\}$, $M_0 + M_1 = I$ then the respective probabilities p_i to observe outcome i are determined by the unknown quantum state via $p_i = \mathrm{Tr}\,(\rho \cdot M_i)$. To fully describe a single qubit we need to measure a set of informationally complete POVMs which will fully define the density matrix. For concreteness let us consider a case when the qubit is represented by the polarization degree of freedom of a single photon. In this case a complete state description can be achieved by estimating the probability of observing one of two orthogonal outcomes in the rectilinear basis (Z, horizontal (h) and vertical (v) polarization), the diagonal basis (X, diagonal (d) and anti-diagonal (a) polarization), or the circular basis (Y, left (l) and right (r) circular polarization). The likelihood of observing a data set $\mathcal{D}$ from these measurements given we know the probabilities of each outcome exactly is

$$P(\mathcal{D}|\alpha) = p_h^{c_h}(1\text{-}p_h)^{c_v} p_d^{c_d}(1\text{-}p_d)^{c_a} p_l^{c_l}(1\text{-}p_l)^{c_r} \tag{4}$$

where $\alpha = \{p_h, p_d, p_l\}$, $\mathcal{D} = \{c_h, c_v, c_d, c_a, c_l, c_r\}$, we have enforced the single basis requirement that the sum of orthogonal probabilities is unity, $p_{h,d,l} + p_{v,a,r} = 1$. Using Bayes rule, the distribution for α given $\mathcal{D}$ is

$$P(\alpha|\mathcal{D}) = \frac{P(\mathcal{D}|\alpha)P(\alpha)}{\int d\alpha P(\mathcal{D}|\alpha)P(\alpha)} \tag{5}$$

which has no quantum constraints, i.e. associated density matrices may not be physical. A physical density matrix ρ for the single-qubit must fulfill constraints

$\mathrm{Tr}\,(\rho) = 1$	probabilities sum to 1
$\langle\phi\|\,\rho\,|\phi\rangle \geq 0$	positive semi-definite
$\rho = \rho^\dagger$	hermitian.

These can all be fulfilled by parametrizing the density matrix as suggested by Daboul [14]. For the single-qubit the parametrized matrix is

$$\rho(\tau) = \begin{pmatrix} \cos^2(u) & \frac{1}{2}\cos(\theta)\sin(2u)\,e^{i\phi} \\ \frac{1}{2}\cos(\theta)\sin(2u)\,e^{-i\phi} & \sin^2(u) \end{pmatrix} \tag{6}$$

where parameter ranges $\tau = \{u, \theta, \phi\}$, $u \in [0, \frac{\pi}{2}]$, $\theta \in [0, \frac{\pi}{2}]$, and $\phi \in [0, 2\pi]$ ensure there is no state redundancy, states having multiple representations. This matrix heeds all

quantum constraints for any values of the parameters. The parameters α in terms of the new parameters τ are

$$p_h(\tau) = \cos^2(u) \tag{7}$$
$$p_d(\tau) = \frac{1}{2} + \frac{1}{2}\sin(2u)\cos(\theta)\cos(\phi) \tag{8}$$
$$p_l(\tau) = \frac{1}{2} + \frac{1}{2}\sin(2u)\cos(\theta)\sin(\phi)\,. \tag{9}$$

This results in the new likelihood

$$P(\mathcal{D}|\tau) = p_h(\tau)^{c_h}(1\text{-}p_h(\tau))^{c_v}p_d(\tau)^{c_d}(1\text{-}p_d(\tau))^{c_a}p_l(\tau)^{c_l}(1\text{-}p_l(\tau))^{c_r}. \tag{10}$$

To complete the new description we must define a new integration measure in τ space. Our original probability space has an infinitesimal length element $(ds)^2 = (dp_h)^2 + (dp_d)^2 + (dp_r)^2$. The measure in this case is reduced to the volume element in Cartesian coordinates $d\alpha = dp_h dp_d dp_l$. This space can be considered a"cube" that includes both physical and unphysical states. Within this cube, the new space is a sphere containing only and all physical density matrices. The length element in this space is [15]

$$(ds)^2 = \mathrm{Tr}\left(d\rho \cdot d\rho^\dagger\right) = \sum_{i,j}\mathrm{Tr}\left(\frac{\partial\rho}{\partial\tau_i}\cdot\frac{\partial\rho}{\partial\tau_j}\right)d\tau_i d\tau_j \tag{11}$$

where $\tau_i \in \{u, \theta, \phi\}$. The new measure, the infinitesimal volume, is

$$d\tau = d\tau_0\; d\tau_1..d\tau_m \mathrm{Det}\sqrt{g} \tag{12}$$

where

$$g_{ij} = \mathrm{Tr}\left(\frac{\partial\rho}{\partial\tau_i}\cdot\frac{\partial\rho}{\partial\tau_j}\right). \tag{13}$$

The integration measure in the ideal single qubit experiment is

$$d\tau = du\; d\theta\; d\phi\; \frac{\sin^3(2u)\sin(2\theta)}{2\sqrt{2}}.$$

As described earlier, this measure is Haar invariant.

We will also consider how this parametrization relates to the Pauli operators

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \tag{14}$$

and their expectations

$$z = \mathrm{Tr}\left(\sigma_z \cdot \rho\right) = \cos(2u) \tag{15}$$
$$x = \mathrm{Tr}\left(\sigma_x \cdot \rho\right) = \sin(2u)\cos(\theta)\cos(\phi) \tag{16}$$
$$y = \mathrm{Tr}\left(\sigma_y \cdot \rho\right) = \sin(2u)\cos(\theta)\sin(\phi). \tag{17}$$

With our likelihood defined, one estimation technique is to approximate the true distribution utilizing Laplace's method [16], also known as the saddle-point approximation. This is a multivariate Gaussian centered on the MLE defined by k parameters. This MLE is found by simultaneously solving k equations of the form

$$\frac{\partial P(\mathcal{D}|\tau)}{\partial \tau_i} = 0 \tag{18}$$

and verifying this point represents the global maximum. The uncertainty in the parameters can be captured utilizing the covariance matrix which we estimate as

$$A_{ij} = -\frac{\partial^2 \log\left(P(\mathcal{D}|\tau)\right)}{\partial \tau_i \partial \tau_j}\bigg|_{\tau=\tau_{\mathrm{ml}}}. \tag{19}$$

The approximate distribution is then

$$P(\mathcal{D}|\tau) \approx \sqrt{\frac{(2\pi)^k}{\det\left(\mathbf{A}\right)}} e^{-\frac{1}{2}(\tau-\tau_{mle})^T \cdot A \cdot (\tau-\tau_{mle})} \tag{20}$$

where τ is a column vector.

For the ideal single-qubit we find unbounded MLE

$$u_{\mathrm{uml}} = \frac{\arccos\left(z_f\right)}{2} \qquad \theta_{\mathrm{uml}} = \arccos\left(\sqrt{\frac{x_f^2 + y_f^2}{1 - z_f^2}}\right) \qquad \phi_{\mathrm{uml}} = \arctan\left(x_f, y_f\right) \tag{21}$$

where z_f, x_f, and y_f are the frequency based linear inversion estimates (LIE) of the Pauli operator expectations

$$z_f = \frac{c_h - c_v}{c_h + c_v} \qquad x_f = \frac{c_d - c_a}{c_d + c_a} \qquad y_f = \frac{c_l - c_r}{c_l + c_r}. \tag{22}$$

When $x_f^2 + y_f^2 + z_f^2 \leq 1$ these LIE are the correct MLE. However the parameter set given in Eq. 21 and 22 is undefined for unphysical states, when $x_f^2 + y_f^2 + z_f^2 > 1$. When this is the case, the MLE is found on the boundary of the Bloch sphere due to the concavity of the likelihood given by Eq. 10. This point is not necessarily the one with smallest Euclidean distance to the unbounded MLE. Determination of the boundary MLE is accomplished by setting $\theta = 0$, restricting us to the boundary, and maximizing the parametrized likelihood Eq. 10 over the parameter ranges $u \in \left[0, \frac{\pi}{2}\right]$ and $\phi \in [0, 2\pi]$. Next, we derive a closed-form BME which always results in a quantum bound obedient estimate.

To calculate the BME for the single qubit density matrix we first evaluate the normalizing constant

$$P(\mathcal{D}) = \int d\tau P(\mathcal{D}|\tau) P(\tau) \tag{23}$$

and then estimate our mean density matrix

$$
\begin{aligned}
\overline{\rho} &= \frac{1}{P(\mathcal{D})}\int d\tau P(\mathcal{D}|\tau)P(\tau)\times\rho\left(\tau\right)\\
&= \frac{1}{P(\mathcal{D})}\int du\; d\theta\; d\phi\;\frac{\sin^3\left(2u\right)\sin\left(2\theta\right)}{2\sqrt{2}}\left(\cos^2 u\right)^{c_h}\left(\sin^2 u\right)^{c_v}\\
&\quad\times\left(\frac{1}{2}+\frac{1}{2}\sin\left(2u\right)\cos\left(\theta\right)\cos\left(\phi\right)\right)^{c_d}\left(\frac{1}{2}-\frac{1}{2}\sin\left(2u\right)\cos\left(\theta\right)\cos\left(\phi\right)\right)^{c_a}\\
&\quad\times\left(\frac{1}{2}+\frac{1}{2}\sin\left(2u\right)\cos\left(\theta\right)\sin\left(\phi\right)\right)^{c_l}\left(\frac{1}{2}-\frac{1}{2}\sin\left(2u\right)\cos\left(\theta\right)\sin\left(\phi\right)\right)^{c_r}\\
&\quad\times\begin{pmatrix}\cos^2\left(u\right) & \frac{1}{2}\cos\left(\theta\right)\sin\left(2u\right)e^{i\phi}\\ \frac{1}{2}\cos\left(\theta\right)\sin\left(2u\right)e^{-i\phi} & \sin^2\left(u\right)\end{pmatrix}.
\end{aligned}
$$

Using the binomial theorem we can rewrite this as

$$
\begin{aligned}
\overline{\rho} \;&= \frac{1}{P(\mathcal{D})}\int_0^{\pi/2}du\int_0^{\pi/2}d\theta\int_0^{2\pi}d\phi\;\frac{8\sin^3\left(u\right)\cos^3\left(u\right)\sin\left(\theta\right)\cos\left(\theta\right)}{\sqrt{2}}\left(\cos^2 u\right)^{c_h}\left(\sin^2 u\right)^{c_v}\\
&\times\sum_{k_d=0}^{c_d}\binom{c_d}{k_d}2^{-k_d}(\sin(u)\cos(u)\cos(\theta)\cos(\phi))^{c_d-k_d}\sum_{k_a=0}^{c_a}\binom{c_a}{k_a}2^{\text{-}k_a}(\text{-}\sin(u)\cos(u)\cos(\theta)\cos(\phi))^{c_a-k_a}\\
&\times\sum_{k_l=0}^{c_l}\binom{c_l}{k_l}2^{-k_r}(\sin(u)\cos(u)\cos(\theta)\sin(\phi))^{c_l-k_l}\sum_{k_r=0}^{c_r}\binom{c_r}{k_r}2^{\text{-}k_l}(\text{-}\sin(u)\cos(u)\cos(\theta)\sin(\phi))^{c_r-k_r}\\
&\quad\times\begin{pmatrix}\cos^2(u) & \frac{1}{2}\cos(\theta)\sin(2u)\,e^{i\phi}\\ \frac{1}{2}\cos(\theta)\sin(2u)\,e^{-i\phi} & \sin^2(u)\end{pmatrix}.
\end{aligned}
$$

The integral over u has solution

$$
\int_0^{\pi/2}du\sin^x u\cos^y u=\frac{1}{2}\,\text{Beta}\left(\frac{1+x}{2},\frac{1+y}{2}\right)\tag{24}
$$

and similar for θ. The integral over ϕ can be shown to be

$$
\int_0^{2\pi}d\phi\;\sin^x\phi\cos^y\phi=\frac{\left(1+(\text{-}1)^x\right)\left(1+(\text{-}1)^y\right)}{2}\,\text{Beta}\left(\frac{1+x}{2},\frac{1+y}{2}\right)\tag{25}
$$

which is zero for odd x or y. To ease representation of the solutions, define

$$
\begin{aligned}
&F_{u_0,u_1,\theta_0,\theta_1,\phi_0,\phi_1}\\
&=\int_0^{\pi/2}du\int_0^{\pi/2}d\theta\int_0^{2\pi}d\phi\;\frac{8\sin^3\left(u\right)\cos^3\left(u\right)\sin\left(\theta\right)\cos\left(\theta\right)}{\sqrt{2}}\left(\cos^2 u\right)^{c_h}\left(\sin^2 u\right)^{c_v}\\
&\quad\times\sum_{k_d=0}^{c_d}\binom{c_d}{k_d}2^{-k_d}(\sin(u)\cos(u)\cos(\theta)\cos(\phi))^{c_d-k_d}\sum_{k_a=0}^{c_a}\binom{n_d\text{-}c_d}{k_a}2^{-k_a}(-\sin(u)\cos(u)\cos(\theta)\cos(\phi))^{c_a-k_a}\\
&\quad\times\sum_{k_l=0}^{c_l}\binom{c_r}{k_r}2^{-k_r}(\sin(u)\cos(u)\cos(\theta)\sin(\phi))^{c_l-k_l}\sum_{k_r=0}^{c_r}\binom{n_c\text{-}c_r}{k_l}2^{-k_l}(-\sin(u)\cos(u)\cos(\theta)\sin(\phi))^{c_r-k_r}\\
&\quad\times\cos^{u_0}(u)\sin^{u_1}(u)\cos^{\theta_0}(\theta)\sin^{\theta_1}(\theta)\cos^{\phi_0}(\phi)\sin^{\phi_1}(\phi)\\
&=\sum_{k_d=0}^{c_d}\sum_{k_d=0}^{c_d}\sum_{k_l=0}^{c_l}\sum_{k_r=0}^{c_r}\binom{c_d}{k_d}\binom{c_a}{k_a}\binom{c_l}{k_l}\binom{c_r}{k_r}2^{-k_d-k_a-k_l-k_r}(\text{-}1)^{c_a+c_r-k_a-k_r}\left(1+(\text{-}1)^{n_c\text{-}k_l\text{-}k_r+\phi_0}\right)\left(1+(\text{-}1)^{n_d\text{-}k_d\text{-}k_a+\phi_1}\right)\\
&\quad\times\text{Beta}\left(\frac{4+2\,c_h+n_d+n_c-k_d-k_a-k_l-k_r+u_0}{2},\frac{4+2\,c_v+n_d+n_c-k_d-k_a-k_l-k_r+u_1}{2}\right)\\
&\quad\times\text{Beta}\left(\frac{2+\theta_0}{2},\frac{1+n_d+n_c-kd-ka-k_l-k_r+\theta_1}{2}\right)\;\text{Beta}\left(\frac{1+n_c-k_l-k_r+\phi_0}{2},\frac{1+n_d-k_d-k_a+\phi_1}{2}\right).
\end{aligned}
$$

The BME for our ideal single-qubit is

$$\overline{\rho} = \frac{1}{F_{0,0,0,0,0,0}} \begin{pmatrix} F_{2,0,0,0,0,0} & F_{1,1,0,1,0,1}+iF_{1,1,0,1,1,0} \\ F_{1,1,0,1,0,1}-iF_{1,1,0,1,1,0} & F_{0,2,0,0,0,0} \end{pmatrix}.$$

This is the best possible estimation of the ideal single-qubit given a set of data $\mathcal{D}$ and a uniform prior.

To illustrate the physicality of the distribution $P(\tau|D)$ and the differences between the MLE, LIE, and the BME consider a true quantum state ρ_0 defined by the parameters $u_0 = 0.864$, $\theta_0 = 0.393$, and $\phi_0 = 5.18$. For visualization we use the Bloch sphere where the state is represented by the expectations of the the Pauli operators given in Eq. 15-17. We simulated taking 10 measurements in the Z, X, and Y bases from which we generated counts $c_h = 7$, $c_v = 3$, $c_d = 7$, $c_a = 3$, $c_l = 0$, and $c_r = 10$. We plot

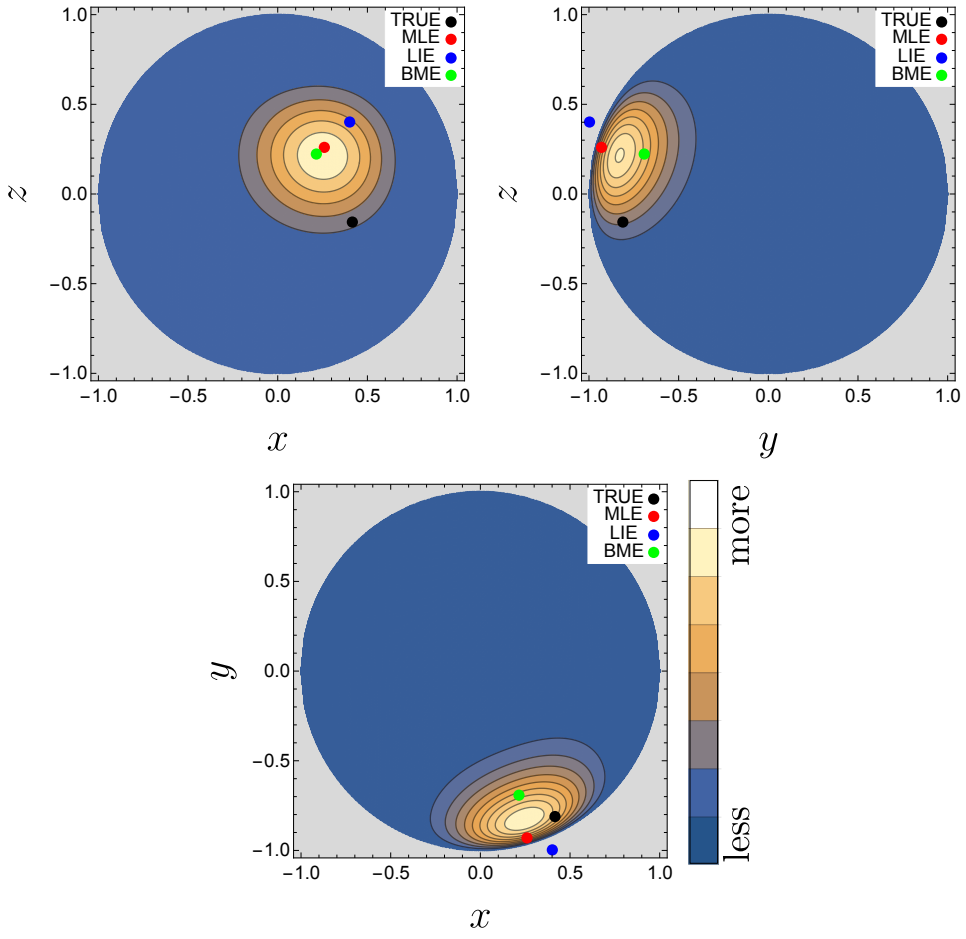

Figure 2. We plot posterior marginal distributions $P(x, z|\mathcal{D})$, $P(y, z|\mathcal{D})$, and $P(x, y|\mathcal{D})$ at top left, top right, and bottom, respectively. The contour plots relate the relative probability of Bloch sphere coordinates. The plotted dots represent the locations of the true state, MLE, LIE, and BME. These plots are given to emphasize the physicality of the posterior distribution used to calculate the BME–the posterior distributions conform to quantum bounds.

	z	x	y	$\sqrt{z^2+x^2+y^2}$
true	-0.156	0.414	-0.813	0.925
MLE	0.263	0.263	-0.928	1.00
LIE	0.400	0.400	-1.00	1.15†
BME	0.226	0.216	-0.695	0.762

Table 1. Bloch sphere coordinates for the true state, MLE, LIE, and BME. †In this case, the LIE is unphysical.

the distributions $P(x, z|\mathcal{D})$, $P(y, z|\mathcal{D})$, and $P(x, y|\mathcal{D})$ in Fig. 2. The coordinates for each estimate are given in Table 1. This small data set emphasizes the parametrized distribution's physicality and the qualitative difference between the MLE and BME. The gray locations correspond to unphysical states. As can be seen in the top right and bottom plots in Fig. 2, the LIE can be unphysical. To correct this, the MLE is found on the boundary, a pure state. In contrast, the BME will always be located within the physical space. This illustration is not made to emphasize the performance of any of specific approach. Performance is addressed in Section 5.

In order to utilize the ideal single qubit formalism with single-qubit experiments, the data can be renormalized to the lowest efficiency measurement similar to the procedure used in Appendix A. This method does not fully utilize the available information and has the additional complication that preliminary experiments must transpire to determine the measurement efficiencies. In the remainder of our manuscript we address qubit estimation for experiments.

4. Bayesian mean estimation for multi-qubit experiments

In an experiment the probability of observing an outcome depends not only on the quantum state but also on the measurement apparatus itself. In this case imperfections and asymmetries in the measurement process prohibit the type of "perfect" estimate we investigated in the last section. Our experiment of investigation is the common two-photon experiment for which we introduce the fundamental assumptions and model below. James et al. [17] previously reported an MLE approach to this experiment as well as higher dimensional experiments. In contrast to that method, we account for qubit loss within our defined likelihood and enable determination of the BME, the best estimate on average [2], which avoids MLE pitfalls such as "zero" probabilities, impossible outcomes.

4.1. Estimating parameters in a single-basis two-photon experiment

In this section, we give an example of estimating parameters in a single-basis experiment. To begin, we assume the existence of a photon pair. A member of this pair is sent to Alice and the other one to Bob each of whom has chosen a measurement basis as seen in Fig. 3. A single photon can result in one of two observable orthogonal outcomes, 0

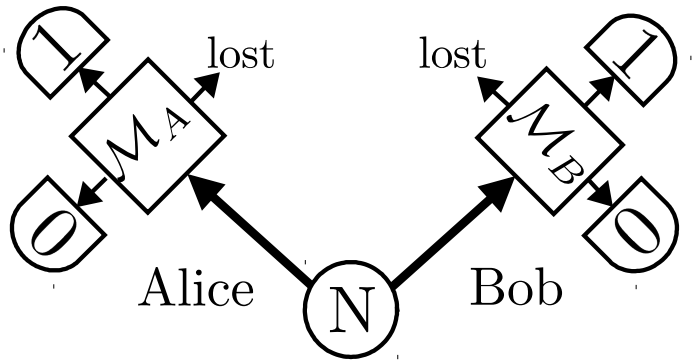

Figure 3. The two-photon experiment is illustrated above. One member of a photon pair, a qubit, is sent to both Alice and Bob who have each chosen a measurement basis. An individual qubit can result in one of two orthogonal outcomes, 0 or 1, or the qubit can be lost.

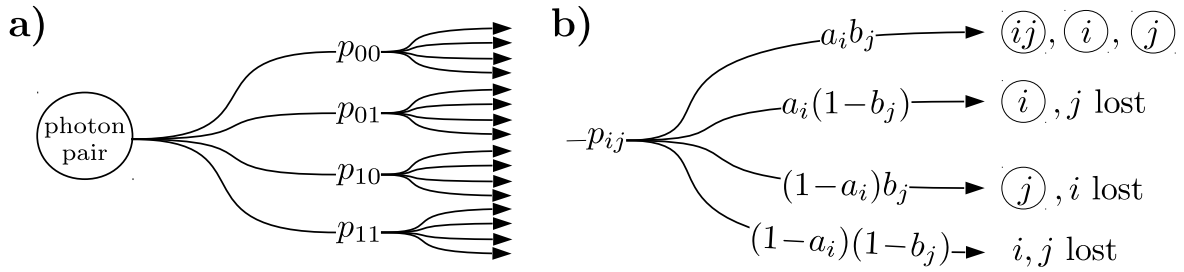

Figure 4. a) Our Bayesian tree begins with the existence of a photon pair. This pair is then "destined" to the joint outcome ij according to probability p_{ij}. b) A closer view of each tree branch shows that each has four possible terminations due to pathway inefficiencies. These possibilities include a joint event, a single event (one photon is lost), and no event (both photons are lost).

or 1, and one unobservable outcome, the photon is lost. All observable outcomes have probabilities of occurrence proportional to the joint probabilities p_{00}, p_{01}, p_{10}, and p_{11} as seen in Fig. 4a. Additionally, Fig. 4b illustrates the four possible outcomes for a given "destiny" when pathway efficiencies are considered. The possibilities include both photons being counted giving one coincidence count and two singles counts, one photon being counted and one lost giving one singles count, or both photons being lost giving no counts.

Alice and Bob record event 0 or event 1 with number $A_0, A_1 \leq N$ and $B_0, B_1 \leq N$, respectively, since typically some portion of the N photons are lost. Losses are due to Alice and Bob's suboptimal pathway efficiencies $\{a_0, a_1, b_0, b_1\} \in [0, 1]$. In the event both members of a photon pair are detected, Alice and Bob observe joint results, giving coincidence totals c_{00}, c_{01}, c_{10}, and c_{11}.

From this data we may enumerate the number of each type of event. The number of joint coincidence events are straightforward, given by the c_{ij} with the probability of these events being $a_i b_j p_{ij}$. The number of events where Alice registers result i and Bob loses

his photon is A_i-c_{i0}-c_{i1}. The probability of this occurrence is $a_i\left[(1-b_0)p_{i0}+(1-b_1)p_{i1}\right]$. The terms for Bob registering a photon and Alice losing her photon are similar. The number of events where both photons are lost is N-A_0-A_1-B_0-B_1+c_{00}+c_{01}+c_{10}+c_{11} with probability

$$p_{\substack{pair\\lost}} = (1\text{-}a_0)(1\text{-}b_0)p_{00} + (1\text{-}a_0)(1\text{-}b_1)p_{01} + (1\text{-}a_1)(1\text{-}b_0)p_{10} + (1\text{-}a_1)(1\text{-}b_1)p_{11}.$$

For now, assume the photon number N is known. In this case, using Bayes' rule, the event number, and the probabilities given above, the PD is

$$P\left(\alpha|\mathcal{D},N\right) = \frac{P\left(\mathcal{D},N|\alpha\right)P(\alpha)}{P\left(\mathcal{D},N\right)} \tag{26}$$

where α=$\{p_{00},p_{01},p_{10},p_{11},a_0,a_1,b_0,b_1\}$ are the unknown parameters, joint probabilities and pathway efficiencies, the data set $\mathcal{D}$=$\{c_{00},c_{01},c_{10},c_{11},A_0,A_1,B_0,B_1\}$ consists of the known singles and coincidence count values totalling

$$s = A_0+A_1+B_0+B_1 \qquad\qquad n = c_{00}+c_{01}+c_{10}+c_{11},$$

respectively,

$$\begin{aligned}
&P(\mathcal{D},N|\alpha) = \\
&\gamma(N)\left(a_0b_0p_{00}\right)^{c_{00}}\left(a_0b_1p_{01}\right)^{c_{01}}\left(a_1b_0p_{10}\right)^{c_{10}}\left(a_1b_1p_{11}\right)^{c_{11}} \\
&\quad\times\left[a_0\left(p_{00}(1-b_0)+p_{01}(1-b_1)\right)\right]^{A0-c_{00}-c_{01}}\left[a_1\left(p_{10}(1-b_0)+p_{11}(1-b_1)\right)\right]^{A1-c_{10}-c_{11}} \\
&\quad\times\left[b_0\left(p_{00}(1-a_0)+p_{10}(1-a_1)\right)\right]^{B0-c_{00}-c_{10}}\left[b_1\left(p_{01}(1-a_0)+p_{11}(1-a_1)\right)\right]^{B1-c_{01}-c_{11}} \\
&\quad\times\left[p_{00}(1\text{-}a_0)(1\text{-}b_0)+p_{01}(1\text{-}a_0)(1\text{-}b_1)+p_{10}(1\text{-}a_1)(1\text{-}b_0)+p_{11}(1\text{-}a_1)(1\text{-}b_1)\right]^{N-(s-n)}, \\
&P(\alpha) = 1, \\
&P(\mathcal{D},N) = \int d\alpha\; P(\mathcal{D},N|\alpha)P(\alpha), \\
&\gamma(N) = \frac{N!}{(N-(s-n))!(A_0\text{-}c_{00}\text{-}c_{01})!(A_1\text{-}c_{10}\text{-}c_{11})!(B_0\text{-}c_{00}\text{-}c_{10})!(B_1\text{-}c_{01}\text{-}c_{11})!c_{00}!c_{01}!c_{10}!c_{11}!}.
\end{aligned}$$

The likelihood $P(\mathcal{D},N|\alpha)$ consists of the probability of each type of event with a multiplicity equal to the number of times it occurred. Both the probabilities and number of events were described in the preceding paragraph. We have retained the full form of the likelihood that includes $\gamma(N)$ for use below.

It is typical in two-photon experiments that the photon number N is not known. If N is known, the following step may be skipped and the above PD is the appropriate choice. Otherwise, we must make N an unobserved parameter or seek a way to eliminate it. Fortunately, there is an analytical method to remove N from the PD completely [5] by taking an average over the N distribution using the summation formula

$$\sum_{m=0}^{\infty}\binom{m+y}{m}m^z x^m = \left(x\frac{d}{dx}\right)^z(1-x)^{-(y+1)}. \tag{27}$$

Since the average is taken over the distribution, see Fig. 5, only probable values of N will have appreciable contribution. Applying this formula, N is removed giving PD

$$P\left(\alpha|\mathcal{D}\right)=\frac{\sum_{N=s-n}^{\infty}P\left(\mathcal{D},N|\alpha\right)P(\alpha)}{P\left(\mathcal{D}\right)}=\frac{P\left(\mathcal{D}|\alpha\right)P(\alpha)}{P\left(\mathcal{D}\right)} \tag{28}$$

where

$$\begin{aligned}
P(\mathcal{D}|\alpha) &= a_0^{A_0}a_1^{A_1}b_0^{B_0}b_1^{B_0}p_{00}^{c_{00}}p_{01}^{c_{01}}p_{10}^{c_{10}}p_{11}^{c_{11}}\\
&\times[p_{00}(1-b_0)+p_{01}(1-b_1)]^{A0-c_{00}-c_{01}}[p_{10}(1-b_0)+p_{11}(1-b_1)]^{A1-c_{10}-c_{11}}\\
&\times[p_{00}(1-a_0)+p_{10}(1-a_1)]^{B0-c_{00}-c_{10}}[p_{01}(1-a_0)+p_{11}(1-a_1)]^{B1-c_{01}-c_{11}}\\
&\times[1\text{-}p_{00}(1\text{-}a_0)(1\text{-}b_0)\text{-}p_{01}(1\text{-}a_0)(1\text{-}b_1)\text{-}p_{10}(1\text{-}a_1)(1\text{-}b_0)\text{-}p_{11}(1\text{-}a_1)(1\text{-}b_1)]^{-s+n-1},\\
P(\alpha) &= 1,\\
P(\mathcal{D}) &= \int d\alpha\ P(\mathcal{D}|\alpha)P(\alpha),
\end{aligned} \tag{29}$$

and, in this specific case,

$$\int d\alpha\equiv\int_0^1 da_0\int_0^1 da_1\int_0^1 db_0\int_0^1 db_1\int_0^1 dp_{00}\int_0^{1-p_{00}} dp_{01}\int_0^{1-p_{00}-p_{01}} dp_{10}$$

with $p_{11}=1-p_{00}-p_{01}-p_{10}$. We omitted all constants.

Assuming the integral can be carried out, we can make estimates of any parameter via its the mean value, for instance,

$$\overline{p_{00}}=\int d\alpha P\left(\alpha|\mathcal{D}\right)\times p_{00}. \tag{30}$$

Likewise, any other parameter mean $\overline{p_{ij}}$, $\overline{a_i}$, or $\overline{b_i}$ as well as their standard deviations may be estimated. One exception is the mean value $\overline{N}$. We find this mean by setting $z=1$ in Eq. (27),

$$\overline{N}=\int d\alpha P\left(\alpha|\mathcal{D}\right)\times\frac{s-n+g(\alpha)}{1-g(\alpha)} \tag{31}$$

where

$$g(\alpha)=p_{00}(1\text{-}a_0)(1\text{-}b_0)+p_{01}(1\text{-}a_0)(1\text{-}b_1)+p_{10}(1\text{-}a_1)(1\text{-}b_0)+p_{11}(1\text{-}a_1)(1\text{-}b_1). \tag{32}$$

In principal, all of the above integrals have analytical solutions via the multinomial theorem,

$$(x_0+x_1+...+x_m)^n=\sum_{k_0+k_1+...+k_m=n}\binom{n}{k_0,k_1,..,k_n}x_0^{k_0}x_1^{k_1}\cdots x_m^{k_m},$$

which gives exact answers in the form of sums of Beta and Gamma functions. However, the computation needed to carry out the resultant sums is prohibitive.

If we cannot efficiently make our parameter estimations analytically, we can utilize numerical sampling to approximate the BMEs of interest. We discuss this in detail in

Appendix B. If the probability of obtaining a sample $\alpha^{(r)}$ tends to the true probability $P(\alpha^{(r)}|\mathcal{D})$, the mean estimations can then be made by repetitive sampling,

$$\overline{\alpha} = \frac{1}{R}\sum_{r=0}^{R}\alpha^{(r)} = \frac{1}{R}\sum_{r=0}^{R}\left\{p_{00}^{(r)}, p_{01}^{(r)}, p_{10}^{(r)}, p_{11}^{(r)}, a_0^{(r)}, a_1^{(r)}, b_0^{(r)}, b_1^{(r)}\right\}$$
$$= \left\{\overline{p_{00}},\ \overline{p_{01}},\ \overline{p_{10}},\ \overline{p_{11}},\ \overline{a_0},\ \overline{a_1},\ \overline{b_0},\ \overline{b_1}\right\}.$$

4.2. Single-basis simulation

Consider a single-basis simulation where unbeknownst to Alice and Bob a source generates N=10,000 photon pairs with joint probabilities and pathways efficiencies

$$p_{00} = 0.3 \qquad p_{01} = 0.05 \qquad p_{10} = 0.2 \qquad p_{11} = 0.45$$
$$a_0 = 0.3 \qquad a_1 = 0.7 \qquad b_0 = 0.9 \qquad b_1 = 0.5\ .$$

The only information available to Alice and Bob are their count numbers

$$A_0 = 1079 \qquad A_1 = 4553 \qquad B_0 = 4474 \qquad B_1 = 2565$$
$$c_{00} = 829 \qquad c_{01} = 89 \qquad c_{10} = 1245 \qquad c_{11} = 1624.$$

Numerical sampling, see Appendix B, is used to produce sample $\alpha^{(r)}$ from $P(\alpha|\mathcal{D})$. Fig. 5 includes histograms for each parameter from 25,200 α samples. Each parameter histogram contains 100 bins. This large sample size was chosen to illustrate that the samples do come from a distribution. For the typical application a much smaller sample size would likely be adequate. From the distribution $P(\alpha|\mathcal{D})$ parameter mean values are found to be

$$\overline{p_{00}} = 0.300 \pm 0.008 \quad \overline{p_{01}} = 0.059 \pm 0.007 \quad \overline{p_{10}} = 0.191 \pm 0.007 \quad \overline{p_{11}} = 0.450 \pm 0.009$$
$$\overline{a_0} = 0.303 \pm 0.011 \quad \overline{a_1} = 0.716 \pm 0.014 \quad \overline{b_0} = 0.918 \pm 0.018 \quad \overline{b_1} = 0.508 \pm 0.011$$

and mean photon number $\overline{N} = 9926 \pm 50.9$. Comparison with the above true values shows qualitative agreement.

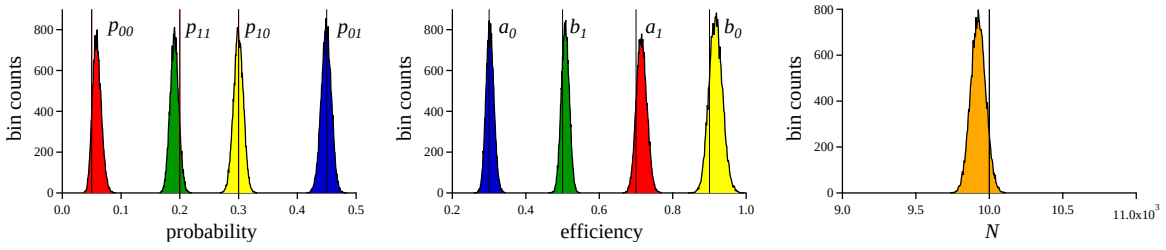

Figure 5. Left. Sample histograms given proportional to the approximate the probability distribution for each parameter p_{ij} whose true value is given by the black vertical line. Middle. Sample histograms are given proportional to the probability distribution for each efficiency parameter a_0, a_1, b_0, and b_1 whose true value is given by the black vertical line. Right. A histogram of samples for parameter N are given proportional to the probability distribution for photon number N.

4.3. Parametrizing the n-dimensional density matrix

For approachability, we obscured the construction of the density matrix for the ideal single qubit given in Eq. 6. We briefly describe this construction here, for a full proof with discussion see [14]. We note that this construction is similar to that recently proposed by Seah et al. [18] whose density matrix sampling application is similar to our approach. Daboul's parametrization can be extended to quantum systems of any dimension. For an n-dimensional Hilbert space, the density matrix is formed using the Cholesky decomposition, requiring $\rho = L(\tau)L(\tau)^\dagger$ with

$$L(\tau) = \begin{pmatrix} L_{11}(\tau) & 0 & 0 & \cdots & 0 \\ L_{21}(\tau) & L_{22}(\tau) & 0 & \cdots & 0 \\ L_{31}(\tau) & L_{32}(\tau) & L_{33}(\tau) & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ L_{n1}(\tau) & L_{n2}(\tau) & L_{n3}(\tau) & \cdots & L_{nn}(\tau) \end{pmatrix}$$

being a lower triangular matrix with positive real diagonal elements. The parameter set τ include n^2-1 parameters which describe a unique density matrix.

The elements L_{ij} may be written as

$$\begin{aligned} L_{ij}(\tau) &= U_i V_{ij} \qquad (j \leq i) \\ L_{ij}(\tau) &= 0 \qquad\quad\ (j > i) \end{aligned}$$

where

$$\begin{aligned} U_1 &= \cos(u_1) & V_{ii} &= 1 \\ U_k &= \cos(u_k) \prod_{j=1}^{k-1} \sin(u_j) \quad (1 < k < n) & V_{i1} &= \cos(\theta_{i1})\, e^{i\phi_{i1}} \quad (i > 1) \\ U_n &= \prod_{j=1}^{n-1} \sin(u_j) & V_{ik} &= \cos(\theta_{ik})\, e^{i\phi_{ik}} \prod_{j=1}^{k-1} \sin(\theta_{ij}) \quad (1 < k < i). \end{aligned}$$

Consider the case of two qubits with dimension n=4, the parametrized matrix

elements of $L(\tau)$ are

$$
\begin{aligned}
L_{11}(\tau) &= \cos(u_1) \\
L_{21}(\tau) &= \sin(u_1)\cos(u_2)\cos(\theta_{21})e^{i\phi_{21}} \\
L_{22}(\tau) &= \sin(u_1)\cos(u_2)\sin(\theta_{21}) \\
L_{31}(\tau) &= \sin(u_1)\sin(u_2)\cos(u_3)\cos(\theta_{31})e^{i\phi_{31}} \\
L_{32}(\tau) &= \sin(u_1)\sin(u_2)\cos(u_3)\sin(\theta_{31})\cos(\theta_{32})e^{i\phi_{32}} \\
L_{33}(\tau) &= \sin(u_1)\sin(u_2)\cos(u_3)\sin(\theta_{31})\sin(\theta_{32}) \\
L_{41}(\tau) &= \sin(u_1)\sin(u_2)\sin(u_3)\cos(\theta_{41})e^{i\phi_{41}} \\
L_{42}(\tau) &= \sin(u_1)\sin(u_2)\sin(u_3)\sin(\theta_{41})\cos(\theta_{42})e^{i\phi_{42}} \\
L_{43}(\tau) &= \sin(u_1)\sin(u_2)\sin(u_3)\sin(\theta_{41})\sin(\theta_{42})\cos(\theta_{43})e^{i\phi_{43}} \\
L_{44}(\tau) &= \sin(u_1)\sin(u_2)\sin(u_3)\sin(\theta_{41})\sin(\theta_{42})\sin(\theta_{43})
\end{aligned}
$$

with $u_i \in [0, \frac{\pi}{2}]$, $\theta_{ij} \in [0, \frac{\pi}{2}]$, and $\phi_{ij} \in [0, 2\pi]$. Indeed, one could instead change the u_i and θ_{ij} trigonometric terms to

$$\cos(u_i) \to \sqrt{u_i'} \qquad \sin(u_i) \to \sqrt{1-u_i'} \qquad \cos(\theta_i) \to \sqrt{\theta_i'} \qquad \sin(\theta_i) \to \sqrt{1-\theta_i'}$$

with $u_i' \in [0,1]$, $\theta_{ij}' \in [0,1]$. The complex terms involving ϕ_{ij} remain unchanged. A similar adjustment was used by Chung and Trueman [19].

4.4. Estimating parameters in a multi-basis two-photon experiment

In section 4.1 and 4.3, respectively, we defined our experimental likelihood for the single-basis experiment and detailed the parametrization of any n-dimensional density matrix. To make estimations using data from multi-basis two-photon experiment we will use both of these pieces. Our example will be full-state tomography. Other multiple basis experiments will have similar estimation constructions. In the case that the data set is incomplete, our method will still return an estimate true to both the given data and all quantum constraints.

To complete full-state tomography Alice and Bob each take measurements in bases Z, X, and Y such that all outcomes are observable in each basis combination ZZ, ZX, XZ, ZY, YZ, XX, XY, YX, and YY. Thus, Alice and Bob's data set will include the data from all 9 basis combinations. The likelihood is a product of the single-basis likelihoods, Eq. 29, from each of these basis combinations,

$$P(\mathcal{D}|\alpha) = P(\mathcal{D}_{ZZ}|\alpha_{ZZ})P(\mathcal{D}_{ZX}|\alpha_{ZX})\cdots P(\mathcal{D}_{YY}|\alpha_{YY}) \tag{33}$$

where α includes the probabilities of all measurement outcomes and the four experimental pathway efficiencies which we assume are the same over all bases. However, in the experimental section, Section 6, we do not make this assumtion. Next, we parametrize our density matrix using the hyperspherical parameters described in the previous section,

$$P(\mathcal{D}|\alpha) \to P(\mathcal{D}|\tau). \tag{34}$$

This parametrization comes with a new measure defined by Eq. 11, 12, and 13.

Putting it all together we can make any BME of interest, for instance the mean density matrix

$$\overline{\rho} = \frac{1}{P(\mathcal{D})} \int d\tau P(\mathcal{D}|\tau) P(\tau) \times \rho(\tau) \tag{35}$$

where $P(\mathcal{D}) = \int d\tau P(\mathcal{D}|\tau) P(\tau)$.

If it is not computationally convenient to evaluate the integrals of the type given in Eq. 35, we can utilize numerical sampling. If we can draw samples $\rho^{(r)}$ from the distribution $P(\tau|\mathcal{D})$ we can estimate the BME of our density matrix as

$$\overline{\rho} = \lim_{R\to\infty} \frac{1}{R} \sum_{i=1}^{R} \rho^{(r)}. \tag{36}$$

We address our numerical sampling approach in Appendix B.

4.5. State certainty

When reporting the values of experimental measurements such as the visibility of an interference curve V or the value of the Bell parameter S, it is typical to provide a standard deviation to describe the uncertainty in the parameter, e.g. $V = 0.98 \pm 0.01$ or $S = 2.65 \pm 0.05$. This gives a quantification of the uncertainty in the estimate. When the BME is multi-dimensional the uncertainty can be represented by a covariance matrix [2, 4]

$$\Delta\rho(\tau) = \begin{pmatrix} \Delta\tau_0^2 & \Delta\tau_0\tau_1 & \cdots & \Delta\tau_0\tau_k \\ \Delta\tau_1\tau_0 & \Delta\tau_1^2 & \cdots & \Delta\tau_1\tau_k \\ \vdots & \vdots & \ddots & \vdots \\ \Delta\tau_2\tau_0 & \Delta\tau_2\tau_1 & \cdots & \Delta\tau_k^2 \end{pmatrix} \tag{37}$$

with each element being a covariance,

$$\Delta\tau_i\tau_j = \overline{\tau_i\tau_j} - (\overline{\tau_i})(\overline{\tau_j}) \tag{38}$$

where $\overline{\tau_i}$ is the expectation, mean value, of τ_i. For $i = j$ this is just the usual variance. Here the k parameters include the $n^2 - 1$ parameters needed to define the density matrix as well as any additional experimental parameters such as the efficiencies.

We can also define a single quantity that captures a compact representation of the certainty in the estimation. We use the *trace distance deviation* ΔD, which is the mean trace distance over the distribution with the distribution's mean density matrix $\overline{\rho}$,

$$\Delta D = \int d\tau D\big(\overline{\rho}, \rho(\tau)\big). \tag{39}$$

The trace distance is

$$D(\rho, \sigma) = \mathrm{Tr}\left(\sqrt{(\rho - \sigma)^2}\right) = \frac{1}{2}\sum_i |\lambda_i| \tag{40}$$

where the λ_i are the eigenvalues of $\rho - \sigma$. We approximate ΔD with numerical sampling using the formula

$$\Delta D \approx \frac{1}{R} \sum_{r=1}^{R} D\big(\,\overline{\rho}, \rho^{(r)}\big). \tag{41}$$

When the certainty is high all samples will be close to the mean value giving $\Delta D \to 0$ which can be compared to the typical standard deviation where smaller is better. This is also useful when there is no particular state one wishes to compare the estimations with such as is typically done when reporting the fidelity.

5. BME performance with numerical sampling

To characterize the performance of the presented estimation methods we used the following procedure. For each $N \in \{10, 10^2, 10^3, 10^4, 10^5\}$ the following steps are repeated:

1. A density matrix ρ is sampled from a uniform distribution using a Haar measure in the hypershperical parameter space.
2. A random set of pathway efficiencies a_0, a_1, b_0, and b_1 are chosen from the range $[0, 1]$. These are chosen to be the same across all bases.
3. We simulate a two-photon experiment for N identical states ρ in each of the 9 bases given in Section 4.4, $9N$ total identical states to generate a data set $\mathcal{D}$. Each simulated experiment for a single basis is the same as described by the Bayesian tree in Section 4.1.
4. Using $\mathcal{D}$ we find using a traditional likelihood and the actual randomly chosen pathway efficiencies as described in Appendix A. Also with $\mathcal{D}$, we find the BME and MLE using the experiment-specific likelihood in which the pathway efficiencies are not known as described in Section 4.4. Thus, the traditional MLE has the unfair and unrealistic advantage of knowing the pathway efficiencies exactly.
5. The distance D, Eq. 40, is found between each estimate and the true state ρ.
6. If the experiment-specific BME or MLE is closer to the true state than the traditional MLE, that estimation type has a win tallied.
7. Steps 1.-6. are repeated for 1000 repetitions.
8. The average distance $\overline{D}$ over all 1000 repetitions is found for the traditional MLE approach and the experiment-specific BME and MLE. The total wins versus the traditional MLE are also recorded.

For these simulations the average distance $\overline{D}$ results are given at left in Fig. 6, and the win percentages for the experiment-specific likelihood MLE and BME versus the traditional MLE are given at right in Fig. 6. The MLE was found using an gradient ascent method described in Appendix C. We emphasize that these results are conservative, since we give the traditional MLE process the pathway efficiencies exactly–these would not be known exactly in an experiment.

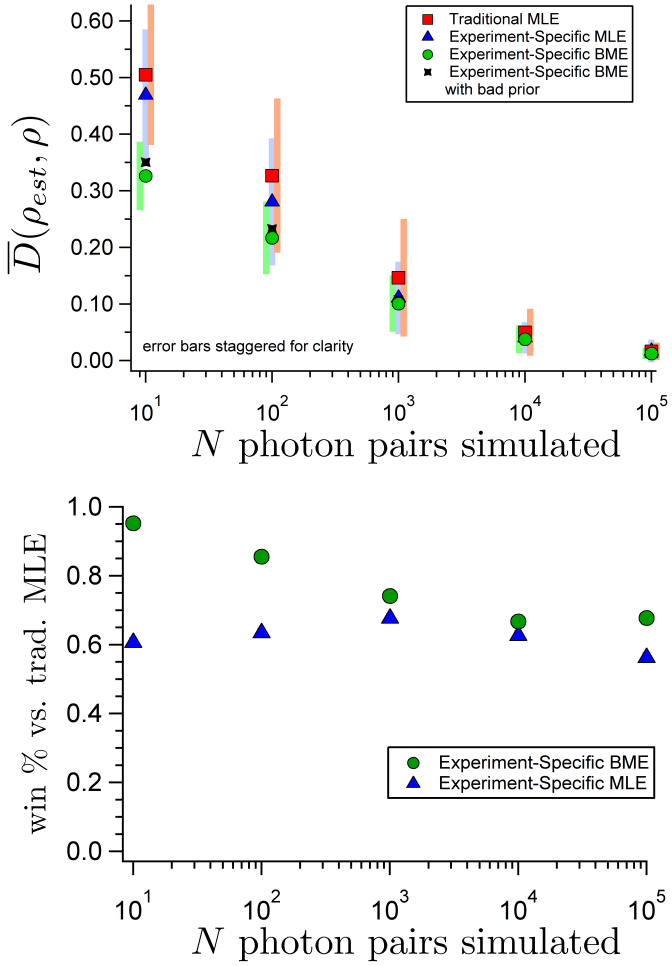

Figure 6. We generated data from simulating two-photon experiments for various states and photon pair number N as outlined in this section. Top. We have plotted the average distance estimate for photon pair number N over 1000 randomly sampled states. Bottom. We give the win percentage for the experiment-specific BME and MLE versus the traditional maximum likelihood method. The best performance is achieved using the experiment-specific Bayesian mean estimate. Another observation from this data is that an experimentalist can achieve a better estimate by switching to an experiment-specific likelihood which allows them to forgo any preliminary experiments to determine normalizing constants.

The states ρ were drawn from a uniform distribution which is Haar invariant when using the measure obtained from Eq. 11, 12, and 13. We also use the same distribution as a prior to compute our BME estimate. We also made estimates using a non-Haar invariant measure to evaluate the significance of prior selection. This results in a drastically different prior relative to that used in generating the random state. In Fig. 6 the experiment-specific BME with the original advantageous prior and the "bad" prior are both plotted. As can be seen, there is, possibly, a small gain using the advantageous prior for the smaller photon pair number estimations. But, it also highlights that the prior choice can be made effectively inconsequential given enough data [5]. The prior

certainly can improve the estimate when little data is available. Granade and colleagues discuss this in depth [4].

6. Experimental tomography

We performed state tomography on a two-photon polarization entangled target state

$$\left|\Psi^+\right\rangle = \frac{1}{\sqrt{2}}\left(|H_A\rangle \otimes |V_B\rangle + |V_A\rangle \otimes |H_B\rangle\right) \tag{42}$$

generated by pumping a periodically poled potassium titanyl phosphate (PPKTP) nonlinear crystal inside a Sagnac loop with two counterpropagating 405nm pump beams [20, 21]. The two possibilities of Type II § spontaneous parametric downconverison (SPDC), either the clockwise or counter-clockwise beam generated a 810nm photon pair, leads to a polarization entangled state output into the idler and signal modes received by Alice and Bob, respectively. Alice and Bob each choose a basis by inclusion or omission of waveplates. Since this requires a physical adjustment to our apparatus for each basis choice, we assume in our likelihood that the efficiencies are independent parameters in each basis. The half-wave and quarter-wave plate matrix operations are, respectively,

$$\mathrm{H} = \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \end{pmatrix} \quad \mathrm{Q} = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}.$$

To measure in basis Z Alice omits her waveplates. To measure in X she includes the half-wave plate, she operates on her single-photon with H. Finally, to measure in the Y basis she includes both waveplates, operating with Q then H. Single-photon detectors record the detection mode, orthogonal outcomes 0 or 1 in each basis. From the experimental data given in Table 1, our mean density matrix is found to be

$$\overline{\rho} = \begin{pmatrix} 0.01 & 0.03+i0.00 & 0.03+i0.00 & -0.00-i0.01 \\ 0.02-i0.00 & 0.48 & 0.48-i0.02 & -0.01-i0.04 \\ 0.03-i0.00 & 0.48+i0.02 & 0.49 & -0.01-i0.05 \\ -0.00+i0.01 & -0.01+i0.04 & -0.01+i0.05 & 0.02 \end{pmatrix}$$

with trace distance deviation $\Delta D = 0.006$ defined in Eq. 39. We have reported only 2 significant digits in $\overline{\rho}$ for brevity. Every element has a finite value, i.e. every outcome has a non-zero probability of occurrence. The fidelity of our mean $\overline{\rho}$ with the intended state Ψ^+ is

$$\mathcal{F} = \sqrt{\langle\Psi^+|\ \overline{\rho}\ |\Psi^+\rangle} = 0.9838 \pm 0.0005.$$

We have not removed accidental coincidences from our estimation, we have assumed this contribution is negligible.

§ The signal and idler photons are produced with orthogonal polarizations in Type II SPDC.

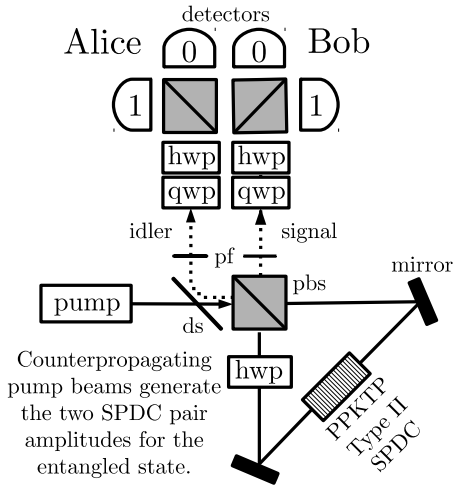

Figure 7. Our two-photon polarization entangled state is generated by pumping a nonlinear PPKTP crystal inside a Sagnac loop with two counterpropagating pump beams each of which may generate Type II SPDC pairs. This leads to a polarization entangled state shared by Alice and Bob. Alice and Bob each choose a basis by inclusion or ommission of waveplates. Single-photon detectors record the detection mode, 0 or 1. ds≡dichroic splitter, pbs≡polarizing beamsplitter, pf≡pump filter, hwp≡ half-wave plate, qwp≡ quarter-wave plate

Basis	A_0	A_1	B_0	B_1	c_{00}	c_{01}	c_{10}	c_{11}
ZZ	47718	50367	45793	44942	189	7302	7903	250
ZX	47117	50726	45467	45831	2735	3826	4075	5061
XZ	45985	51051	45509	44441	4077	3643	3806	4317
ZY	47775	51018	46149	45415	2579	4382	4650	4545
YZ	44564	49626	45739	44157	3382	4155	4414	3505
XX	46547	50920	45186	45658	6801	104	148	9083
XY	45630	50932	44970	44155	3131	3770	3309	4638
YX	44553	49430	45364	45428	3775	3318	2909	4650
YY	44499	49666	45718	45152	6586	61	177	8915
Dark	418	460	406	440	0	0	0	0

Table 2. Experimental tomography data for our two-photon experiment. Counts were 1 second in length. A final dark count was taken with the photon source blocked.

7. Conclusions

We have presented a novel method of Bayesian mean estimation using hyperspherical parametrization and an experiment-specific likelihood. This method has allowed us to derive a closed-form BME for the ideal single-qubit and to develop a numerical approach to approximating the BME for a two-qubit experiment using numerical slice sampling. Our approach offers the real world benefit of eliminating the need for preliminary experiments in common two-photon experiments by accounting for qubit loss within the likelihood. Our method is also scalable beyond two-qubit systems. Finally, we illustrated our approach by applying it to the measurement data obtained from a real-

world two-photon entangled state.

8. Acknowledgement

We would like to thank Nick Peters, Ryan Bennink, and Robin Blume-Kohout for comments, criticisms, and suggestions regarding this manuscript. This work was supported by the Oak Ridge National Laboratory Postdoctoral Program. This manuscript has been authored by UT-Battelle, LLC, under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy.

9. References

[1] Z. Hradil. Quantum-state estimation. *Phys. Rev. A*, 55:R1561–R1564, Mar 1997.

[2] Robin Blume-Kohout. Optimal, reliable estimation of quantum states. *New Journal of Physics*, 12(4):043034, 2010.

[3] KRW Jones. Principles of quantum inference. *Annals of Physics*, 207(1):140–170, 1991.

[4] Christopher Granade, Joshua Combes, and D G Cory. Practical bayesian tomography. *New Journal of Physics*, 18(3):033024, 2016.

[5] Edwin T Jaynes. *Probability theory: the logic of science*. Cambridge university press, 2003.

[6] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

[7] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential monte carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436, 2006.

[8] F. Huszár and N. M. T. Houlsby. Adaptive bayesian quantum tomography. *Phys. Rev. A*, 85:052120, May 2012.

[9] Daniel F. V. James, Paul G. Kwiat, William J. Munro, and Andrew G. White. Measurement of qubits. *Phys. Rev. A*, 64:052312, Oct 2001.

[10] Cyril Stark. Self-consistent tomography of the state-measurement gram matrix. *Physical Review A*, 89(5):052109, 2014.

[11] Robin Blume-Kohout, John King Gamble, Erik Nielsen, Jonathan Mizrahi, Jonathan D Sterk, and Peter Maunz. Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit. *arXiv preprint arXiv:1310.4492*, 2013.

[12] Radford M. Neal. Slice sampling. *The Annals of Statistics*, 31(3):pp. 705–741, 2003.

[13] Radford M. Neal. Markov chain sampling methods for dirichlet process mixture models. *Journal of Computational and Graphical Statistics*, 9(2):249–265, 2000.

[14] Jamil Daboul. Conditions on density matrix elements and their application in resonance production. *Nuclear Physics B*, 4(1):180–188, 1967.

[15] Yan V Fyodorov. Introduction to the random matrix theory: Gaussian unitary ensemble and beyond. *London Mathematical Society Lecture Note Series*, 322:31, 2005.

[16] David JC MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.

[17] Daniel F. V. James, Paul G. Kwiat, William J. Munro, and Andrew G. White. Measurement of qubits. *Phys. Rev. A*, 64:052312, Oct 2001.

[18] Yi-Lin Seah, Jiangwei Shang, Hui Khoon Ng, David John Nott, and Berthold-Georg Englert. Monte carlo sampling from the quantum state space. ii. *New Journal of Physics*, 17(4):043018, 2015.

[19] S. U. Chung and T. L. Trueman. Positivity conditions on the spin density matrix: A simple parametrization. *Phys. Rev. D*, 11:633–646, Feb 1975.

[20] Taehyun Kim, Marco Fiorentino, and Franco N. C. Wong. Phase-stable source of polarization-entangled photons using a polarization sagnac interferometer. *Phys. Rev. A*, 73:012316, Jan 2006.

[21] Brian P. Williams, Keith A. Britt, and Travis S. Humble. Tamper-indicating quantum seal*. *Phys. Rev. Applied*, 5:014001, Jan 2016.

Appendices

A. Maximum likelihood estimation with a traditional likelihood

Traditional MLE relies on a simple multinomial likelihood that relates the probability of an observation to its quantum probability using Born's Rule,

$$P(\mathcal{D}|\alpha) = \prod_i \mathrm{Tr}\left(E_i \cdot \rho\right)^{c_i} \tag{43}$$

where E_i are the observable of interest, for instance the Pauli operators σ_z, σ_x, σ_y or possible Kronecker products of them, $\sigma_z \otimes \sigma_x$, and so forth as dimensionality demands. For instance, the likelihood for a two-photon state with measurement results from a single-basis is

$$P(\mathcal{D}|\alpha) = p_{00}^{c_{00}} p_{01}^{c_{01}} p_{10}^{c_{10}} \left(1 - p_{00} - p_{01} - p_{10}\right)^{c_{11}} \tag{44}$$

with $\mathcal{D} = \{c_{00}, c_{01}, c_{10}, c_{11}\}$ and $\alpha = \{p_{00}, p_{01}, p_{10}\}$. The problem with this simple view is that only in an ideal, or possibly perfectly symmetric, experiment does the data set $\mathcal{D}$ truly result from the quantum probabilities alone. Instead the measurement apparatus will add its own bias and inefficiency to the result. Thus, the data set may need to be corrected using experimental constants determined from initial experiments. In this paper we focus on the pathway efficiencies a_0, a_1, b_0 and b_1 present in the two-photon apparatus. Given that we know these exactly for a single basis we can correct the data set in the following manner. Identify the smallest efficiency for both Alice and Bob

$$a_m = \min(a_0, a_1) \qquad b_m = \min(b_0, b_1) \tag{45}$$

and use these to correct the counts to

$$k_{ij} = \frac{a_m b_m}{a_i b_j} c_{ij}. \tag{46}$$

If a_0=a_1 and b_0=b_1, it should be apparent that no correction is needed. With the corrected data set $\mathcal{D}'$=$\{k_{00}, k_{01}, k_{10}, k_{11}\}$ maximization of $P(\mathcal{D}'|\alpha)$ can proceed.

B. Numerically sampling density matrices

In this Appendix we describe our procedure for sampling density matrices from the true distribution. While other methods of numerical sampling can be used, we have utilized slice sampling [12, 16]. We give a brief description of slice sampling from a single-parameter distribution before describing its extension to density matrix sampling.

B.1. Slice sampling

Slice sampling is an approach to sampling parameter values x from the distribution $P(x)$ given that only the unnormalized distribution $P^*(x)$ is known [12, 16]. This is useful when evaluating $P^*(x)$ everywhere is resource prohibitive–the normalization $Z{=}\int dx P^*(x)$ is unknown. In this case, consider Fig. 8a) where the dashed curve represents the likelihood $P^*(x)$ which we do not know everywhere but can evaluate anywhere. Now consider that we can uniformly fill the space under this curve with points as depicted in Fig. 8b). If we forget the "y" coordinate and randomly select one of the R points we would tend to be sampling from the true distribution as $R \to \infty$. Slice sampling is one method by which to uniformly fill the space under $P^*(x)$ with points.

There are different slice sampling methods, but we will use that introduced by Neal [12, 16]. Below we evaluate the unnormalized distribution $P^*(x)$ in our sampling algorithm. Referring to Fig. 9,

a) Start at point x_0. We assume, for now, this is a random point.

b) At this location a vertical coordinate $y = q \cdot P^*(x)$, $q \in [0, 1]$, is chosen uniformly.

c) From x_0 we "step out" to the left and to the right in steps of w until the left point

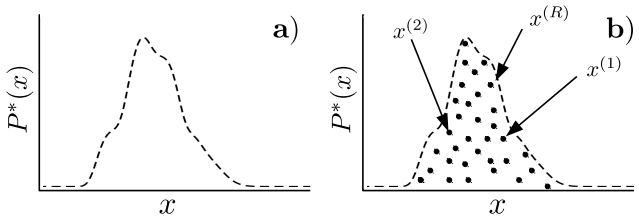

Figure 8. If one could fill the space under the curve of an unknown and unnormalized distribution $P^*(x)$ uniformly, each of these points would represent a sample from the true distribution $P(x)$ as the number of points goes to infinity.

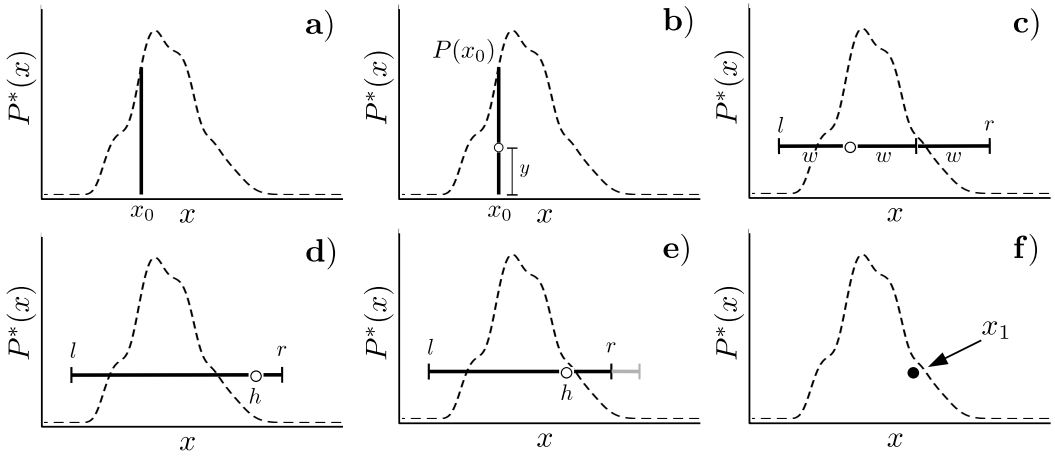

Figure 9. This sequence depicts one loop of the the slice sampling algorithm.

l and right point r are both outside the curve, $P^*(l) < y$ and $P^*(r) < y$.

d) We uniformly sample a new point h from the interval $x \in [l, r]$. If h lies inside the curve, $P^*(h) \geq y$ we accept this as our new point. If h lies outside the curve, $P^*(h) < y$, we set the new leftmost or rightmost point equal to h. Whether h is greater or lesser than x_0 reveals which. We do this for our expected unimodal distribution to increase the chance of acceptance for the next sample by eliminating known rejection regions. We reject the first point and shrink the interval.

e) We select and accept new point $h = x_0$ from the interval, since $P^*(x_1) > y$.

f) We start over at step a) from point x_1.

In reality, we use $\ln P^*(x)$ to compare points in the distribution which is more numerically convenient. The first point x_0 in step a) has not come from the true distribution. To resolve this, samples can be taken and forgotten until samples are being taken from $P(x)$ with no influence of the starting point x_0 present. To evaluate when this point has been reached we utilize multiple independent samplers. Our stopping criteria should ensure that all samplers have converged, and also that enough samples have been taken to sufficiently approximate the distribution $P(x)$. Using multiple samplers has compound utility in that convergence can be assessed and also that more samples can be taken in less time when samplers are run in parallel. Our sampling procedure is to perform a "burn-in" where each sampler takes k_0 samples, the mean $\overline{x}$ and the mean standard deviation σ are calculated for each sampler. When the mean of the ith sampler $\overline{x}_i$ is separated from each of the other j sampler means $\overline{x}_j$ by less than $m\sigma_j$, m smaller requires closer convergence, we stop the burn in. If this criteria is not met after k_i samples, we forget all but the last sample point, double the number of samples $k_{i+1} = 2k_i$, and start over. Once the criteria is met, we repeat the last sampling request for each sampler. The data from this final request is combined as the final set of samples. These samples tend to be from the true distribution $P(x)$ as the convergence parameter $m \to 0$.

The estimations we have made in this article have been from distributions assumed to be unimodal. While this is certainly true for the simplest distributions given, it is not conclusive for the more complex distributions we have derived. However our assumption is corroborated by the simulations where multiple samplers are given the same data set, started in random independent locations, and always converge. However, we have not addressed how to approach known multi-modal distributions. In this case, more advanced methods such as sequential Monte Carlo (SMC) [7] must be applied to full characterize the posterior distribution.

B.2. Density matrix sampling

The goal of our slice sampling application is to sample density matrices from the true but unknown PD $P(\tau|\mathcal{D}) = \frac{1}{Z}P(\mathcal{D}|\tau)P(\tau)$ used in Eq. 36. To do this we slice sample parameter x in τ from the conditional unnormalized distribution $P^*(x|\tau')$ by keeping all other parameters τ' fixed, $x \notin \tau'$. The sampling procedure is performed sequentially

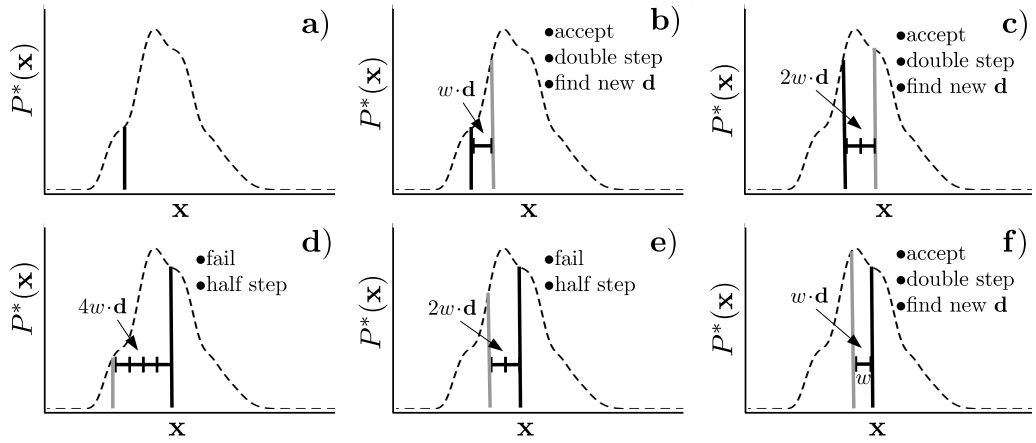

Figure 10. This sequence depicts a few iterations of the gradient ascent method. The bottom axis represents an idealized axis along the line of highest gradient ascent direction within the multi-dimensional space.

for k, $k \geq n^2 - 1$, parameters while keeping each new parameter sample point for the next following parameters' conditional likelihoods. Once each parameter has been sampled the new points constitute one sample $\rho^{(r)}$. Beginning with sample point $\rho(\tau^i)$, $\tau^i = \{\tau_1^i, \tau_2^i, \cdots, \tau_k^i\}$,

1) sample t_1^{i+1} from $P^*(\tau_1 | \tau_2^i, \cdots, \tau_k^i)$
2) sample t_2^{i+1} from $P^*\left(\tau_2 | \tau_1^{i+1}, \cdots, \tau_k^i\right)$
j) sample t_j^{i+1} from $P^*\left(\tau_j | \tau_1^{i+1}, \cdots, \tau_{j-1}^{i+1}, \tau_{j+1}^i, \cdots, \tau_k^i\right)$
k) sample t_k^{i+1} from $P^*\left(\tau_k | \tau_1^{i+1}, \tau_2^{i+1}, \cdots, \tau_{k-1}^{i+1}\right)$.

The $(i+1)^{\text{th}}$ sample is $\rho^{(i+1)} = \rho(\tau^{i+1})$.

Each parameter may have unique ranges, thus slice sampling for each must account for the minimum and maximum parameter values $x \in \{\min, \max\}$. Additionally, some parameters can be cyclic. For instance, angle $\phi \in \{0, 2\pi\}$. Extra consideration must be taken for these parameters since their distribution can be centered on a boundary. Hard bounds would not allow this. For these parameters we gave no bounds during the slice sampling algorithm.

C. Maximization using gradient ascent

To locate a local maximum from the likelihood we use an gradient ascent method as shown in Fig. 10. This methods auto-tunes the step size to make both large and small steps when appropriate.

a) Starting from the current multi-dimensional point $\mathbf{x}$ the gradient $\nabla P^*(\mathbf{x})$ is determined using the finite difference method. The direction of ascent $\mathbf{d}$ is found

by normalizing the gradient with the ℓ^2 norm

$$\mathbf{d} = \frac{\nabla P^*(\mathbf{x})}{|\nabla P^*(\mathbf{x})|}$$
$$|\nabla P^*(\mathbf{x})| = \sqrt{\sum_i \left(\frac{\partial P^*(\mathbf{x})}{\partial x_i}\right)^2}$$

b) A step of size w is made in the direction of increase, $\mathbf{x} \to \mathbf{x} + w \cdot \mathbf{d}$. This point is accepted since $P^*(\mathbf{x} + w \cdot \mathbf{d}) \geq P^*(\mathbf{x})$. The step size is doubled.

c) The direction of ascent $\mathbf{d}$ is found. A step of $2w$ is made in the direction of increase, $\mathbf{x} \to \mathbf{x} + 2w \cdot \mathbf{d}$. This point is accepted since $P^*(\mathbf{x} + 2w \cdot \mathbf{d}) \geq P^*(\mathbf{x})$. The step size is doubled.

d) The direction of ascent $\mathbf{d}$ is found. A step of $4w$ is made in the direction of increase, $\mathbf{x} \to \mathbf{x} + 4w \cdot \mathbf{d}$. This point is not accepted since $P^*(\mathbf{x} + 4w \cdot \mathbf{d}) < P^*(\mathbf{x})$. The step size is halved.

e) A step of $2w$ is made in the direction of increase, $\mathbf{x} \to \mathbf{x} + 2w \cdot \mathbf{d}$. This point is not accepted since $P^*(\mathbf{x} + 2w \cdot \mathbf{d}) < P^*(\mathbf{x})$. The step size is halved.

f) A step of w is made in the direction of increase, $\mathbf{x} \to \mathbf{x} + w \cdot \mathbf{d}$. This point is accepted since $P^*(\mathbf{x} + w \cdot \mathbf{d}) \geq P^*(\mathbf{x})$. The step size is doubled.

In this way, the local maximum is ultimately approached.